

THE WORKSHOP

Prioritisation

A 60-to-90 minute pattern for picking the next order of work when the backlog has outgrown the room's memory. MoSCoW, RICE, Cost of Delay, value-vs-effort – not as rival religions, but as lenses chosen to fit the constraint. Silent scoring, reveal, argue the fringes, land the order.

2026-06-26

barkingiguana.com/writing/the-workshop-prioritisation/

Contents

Prioritisation	3
What's It For	3
What It's Not For	4
Definitions & Background	4
Inputs	5
Outputs	5
Who's Needed	6
How To Run It	7
What Can Go Wrong	10
Next Steps	11
Variants	12

Too many ideas, not enough time, and the loudest stakeholder always wins. The Now/Next/Later workshop turns “everything is urgent” into a plan that connects sprint work to quarterly outcomes. Worked example: [What Changes First](#).

Prioritisation

Prioritisation names the constraint that makes prioritisation necessary, picks a scoring lens that fits, scores silently, reveals together, and leaves the room with a stated order, inside 60 to 90 minutes. Sometimes called backlog grooming at the strategic level, portfolio sequencing, or stack-ranking. Frequently confused with estimation (which sizes work rather than ordering it) and with roadmapping (which lays work across time rather than deciding which piece comes next). The frameworks people associate with the session, MoSCoW (Dai Clegg, 1994), RICE (Intercom, 2014), Cost of Delay / WSJF (Don Reinertsen), value-vs-effort, are tools, not the technique. The technique is the shape of the session.

At a glance

- *Who, for how long:* a facilitator who doesn't score, a product lead with veto, two or three engineers, and a CS / UX / operations voice. Four to seven people, 60 to 90 minutes.
- *What you walk out with:* a stated, committed order of work the product lead reads aloud, a deliberate “not now” list, the hidden context that the silent-scoring distribution surfaced, and a defensible trail (constraint card, framework, scores) for the stakeholders who'll ask later.
- *When to reach for it:* a backlog that's outgrown the room's memory, a constraint about to bite (quarter end, release window, capacity), or stakeholders looping on order. Not for unrefined backlogs, sessions with no agreed constraint, or strategic *what-should-we-do* questions (use [Impact Mapping](#) or [Business Model Canvas](#) first).

What's It For

Two people in a team disagree about what to build next. Neither is wrong. One is carrying the risk of a payment bug they saw last week; the other is carrying the deadline of a marketing campaign they promised in last month's all-hands. They both have reasons. Those reasons never get placed on a table together, so the argument happens in sprint planning, in Slack, in corridor conversations, and the order of work drifts toward whoever argued most recently.

Prioritisation is not about frameworks. The frameworks are lenses. The hard work is naming the constraint that *makes* prioritisation necessary in the first place: a budget, a headcount, a quarter end, a release window, a founder's attention. Without the constraint named, every framework produces a different answer, and the answer people accept is whichever one confirms what they already wanted.

This session exists to force the constraint to the surface, pick a lens that fits, and collect the scoring as data rather than as argument. Silent scoring is the forcing function. When everyone writes their score at the same time without seeing anyone else's, the distribution of scores is itself the signal. If everyone agrees, move on. If everyone disagrees, that's the conversation worth having.

Friction is a feature of this session. If the room leaves without anyone having had to revise their opinion, the session probably didn't produce a prioritisation; it produced a ritual.

Reach for it when:

- A backlog has grown beyond what the team can hold in working memory and decisions are drifting
- A constraint is about to bite: end of quarter, a release window, a budget, a hiring slowdown
- Two or more stakeholders are arguing about work order and the arguments keep looping
- A new initiative is starting and you need to pick the first three items out of twenty
- The team is about to commit and you want disagreement surfaced before commitment, not after

What It's Not For

Skip it when:

- The backlog hasn't been refined. Prioritising work that the team doesn't understand produces theatre.
- The constraint is not yet agreed. A session without a named constraint is a preference poll.
- The work is strategic rather than tactical. Use Impact Mapping or Business Model Canvas to decide *what* before ordering it.
- You don't have the person who can veto. Prioritisation without authority is advisory at best.

Stop a session that's already started if:

- The constraint can't survive scrutiny. If the room can't agree on what they're optimising against in fifteen minutes, the prioritisation isn't the problem; something upstream is
- Scores are bunching because people are conforming rather than thinking
- The product lead won't commit at the end

Stopping and re-framing the constraint is not failure. Running a scoring session against a fake constraint and committing to the output is.

Definitions & Background

This is the one reference section it's worth reading twice. Picking the wrong framework doesn't break the session, but it makes the output harder to use.

Now / Next / Later. A horizon-based roadmap framing. Good when the constraint is *attention* rather than capacity: the team wants to commit confidently to the near term, signal direction for the medium term, and keep options open further out. Now is what the team commits to; Next is the short-list immediately after; Later is the long-list deliberately not committed to. The fuzziness of *Later* is the point. Weak when items inside a horizon need precise ordering; use one of the lenses below for that. Use for: roadmap conversations with stakeholders, anything where pretending to plan twelve months out has historically backfired.

MoSCoW (Must, Should, Could, Won't). A scope-management tool for a release. Good when the constraint is a deadline and the output is a yes/no decision on each item. Weak for long-lived backlogs: everything drifts into Must, and the categories stop separating anything. Use for: release planning, MVP scoping, time-boxed launches.

RICE (Reach × Impact × Confidence / Effort). A growth-backlog tool. Good when items have measurable reach (users, requests, revenue) and the team is choosing between many similar-shape experiments. Weak when the items are qualitatively different; you end up comparing a marketing experiment to a migration to a support tool, and the scores are incomparable. The honest cheat surface: *Confidence is the dial people turn to make their item win*. Cap it at three values (low / medium / high) and force everyone to justify any “high” out loud. Use for: growth experiment backlogs, homogeneous feature pipelines.

Cost of Delay / WSJF: (User value + Time value + Risk reduction) / Job size. A portfolio-sequencing tool. Good when items have different urgency profiles: some things cost the business by the day, others by the quarter, some only once. Surfaces the items whose delay is cheap even when their value looks high. Weak when the team can’t honestly estimate the dollar impact of delay; the scores become theatre. The honest cheat surface: *Job Size is your estimate’s estimate; it inherits all the team’s sizing pathologies*. Don’t let one person produce all four numbers; spread the responsibility. Use for: mixed investment fleets, platform work alongside feature work, “we’re behind on everything” situations.

Value-vs-effort 2×2. A light-weight quadrant. Good for a fortnightly top-up of the backlog or a small team with a short horizon. Weak when the room needs a precise order: the 2×2 gives you quadrants, not a rank. Use for: tactical top-of-backlog work, opportunistic fortnights, any team under six people.

A note on choosing. The framework is a lens on the constraint. Now/Next/Later assumes attention is the limit. MoSCoW assumes a deadline. RICE assumes reach. Cost of Delay assumes time-sensitivity. Value-vs-effort assumes rough is good enough. Read the constraint; pick the lens that matches.

A note on combining. Now/Next/Later and the scoring frameworks aren’t alternatives; they’re complementary. Score with whichever lens fits the constraint, then bucket the scored output into Now / Next / Later for stakeholder-facing roadmap conversations. The buckets carry the commitment signal; the scores carry the justification.

Inputs

- A list of 10 to 30 candidate items, each with a one-sentence description the room recognises.
- A named constraint on a card. Written down before the session, on the wall throughout it.
- A rough sense of the effort involved: not an estimate; a t-shirt size is enough.
- The right people in the room (see *Who’s Needed*), and a 60-to-90 minute slot with no interruptions.

If any of these are missing, the session isn’t ready to run.

Outputs

What lands at the end:

- A stated, committed order of work that survives the week, read aloud by the product lead and confirmed by the room.
- A “not now” list: the items deliberately deprioritised, just as valuable as the top of the order.
- Hidden context surfaced. The 2-versus-9 scoring distributions reveal information one person had and others didn’t.

- A defensible trail for stakeholders who ask “why are we doing this and not that”: the constraint card, the framework chosen, the scores, the order.
- Photographs of the scoring wall, the constraint card, and the final ordered list.

These outputs feed straight into:

- **Sprint Planning**. Prioritisation produces the input to sprint planning. A sprint planning session with unprioritised items is really a prioritisation session in disguise.
- **Impact Mapping**. Impact Mapping decides what work *deserves* to be on the list; prioritisation decides the order of what made it on. Run Impact Mapping first for a new initiative; run prioritisation when the list already exists.
- **User Story Mapping**. Story Mapping slices a journey into release chunks; prioritisation orders the items inside a slice. Composes naturally: Story Mapping feeds prioritisation a shortlist.
- **Assumption Mapping**. When two items score similarly but one is full of assumptions and the other isn't, prioritise the one with evidence. Assumption Mapping is the sanity check before committing to a scored order.
- **Business Model Canvas**. The Canvas sets the strategic shape; prioritisation sequences the work inside it. Run the Canvas first when the prioritisation argument is really a strategy argument in disguise.
- **Jobs to be Done**. JTBD names the jobs the work is serving; prioritisation orders which job to serve first. Without JTBD, scoring often measures excitement; with it, scoring measures service to a named job.

Who's Needed

Four to seven people, 60 to 90 minutes:

- **Facilitator**. Runs the clock, holds the framework, and protects silent scoring from being interrupted. The facilitator does not score; they hold the shape of the room.
- **Product lead**. Mandatory. They own the decision. The session produces input; the product lead commits to the order at the end. Without them, the session is a recommendation, and recommendations rot.
- **Engineers**. Two or three. They carry the effort estimates and the technical risk, and they catch the items whose “small” estimate hides a week of surprise.
- **CS / UX / Operations voice**. Someone who talks to customers or lives with the consequences of the current system. They weight items against lived pain that doesn't show up on dashboards.
- **Founder or senior stakeholder (optional)**. Useful when the constraint involves budget, headcount, or strategic bets. Disruptive when their presence prevents honest disagreement. If they attend, the facilitator has to be willing to redirect them.

Below four and the scoring distribution has no shape; above seven, silent scoring takes twice as long and the argue-the-fringes phase fragments.

Who to leave out:

- Large review boards. Prioritisation by committee is the thing this session is trying to replace.

- People who will be relayed the outcome. They don't need to be in the room; a short written summary works.
- Observers who can't resist commenting. Silent scoring breaks if people are narrating.

How To Run It

Phase	Duration	Materials	Key question
Frame the constraint	10 min	Constraint card	"What are we optimising against?"
Propose the framework	10 min	Framework card	"Which lens fits this decision?"
Silent scoring	15 min	Score sheet per person	"What does each item score?"
Reveal and discuss	25 min	Combined scores on wall	"Where do we agree? Where don't we?"
Argue the fringes	20 min	Top and bottom of the list	"Are the edges correct?"
Land the order and commit	10 min	Final ordered list	"Who owns what next?"
Total	~90 minutes		

Sixty minutes is possible with a smaller backlog and a team that has run the pattern before. Ninety is the realistic number. The silent scoring phase refuses to be compressed; the time is the forcing function.

Three modes run through the session, and the transitions between them matter:

- Named constraint. Written on a card, on the wall, throughout the session. Every argument returns to it. *"We're optimising for launch by end of Q2; that's on the wall."*
- Silent scoring. Each person scores every item alone, on paper or in a private tab. No talking, no eye contact. The facilitator's job is to protect the silence and to resist the urge to clarify items mid-scoring (capture questions, answer them once everyone is done).
- Open discussion. After the reveal. Directed by the scoring distribution, not by whoever speaks first. The facilitator walks the items with the widest disagreement and opens each one deliberately.

The framework is chosen before scoring begins and doesn't change mid-session. Picking a different framework halfway through invalidates the scores and wastes the silent phase.

Phase 1: Frame the constraint (10 min)

Walk to the wall, put the constraint card up, and read it aloud:

"Our constraint for this session is: we have engineering capacity for six stories in the next fortnight and we need to decide which six. That's on the wall. Everything we do for the next eighty minutes is in service of that constraint."

Then check that everyone agrees what the constraint means. Five minutes of “does ‘capacity’ include the on-call rotation?” or “are we counting the half-timer at full capacity?” now is cheaper than an hour of argument later.

What to watch for:

- The imposed constraint. Someone says *“I don’t agree with this constraint.”* Pause. If the constraint is imposed from outside the room, name that: *“I agree this isn’t the constraint we’d have chosen, but it’s the one we’re working inside today. If we want to revisit it, that’s a separate conversation with leadership.”*
- The fuzzy constraint. *“We need to ship something good.”* Not a constraint. Push: *“By when? For whom? At what cost?”* If the answer isn’t concrete, end the session and schedule one to sharpen the constraint.
- Multiple constraints at once. *“End of quarter and keep the on-call load down.”* Pick the dominant one. You can’t optimise for two things without a trade-off rate; name it.

Phase 2: Propose the framework (10 min)

State the framework you’re going to use and why, briefly. Don’t take a vote; the facilitator proposes, the room refines.

“Given the constraint is a fixed capacity and the items are a mix of features, fixes, and a migration, I’m going to propose Cost of Delay. The alternative would be value-vs-effort, which is lighter but won’t separate the migration from the features. Does anyone want to push back?”

Allow challenge. If the room genuinely prefers a different framework and the facilitator agrees it fits, switch. If it’s a matter of preference, stick with the proposal; the goal is scoring, not framework debate.

What to watch for:

- Framework debate eating the session. Park it after five minutes: *“Any framework will give us a useful signal. Let’s score.”*
- A framework that doesn’t match the constraint. MoSCoW for a long-lived portfolio will mis-categorise. Redirect: *“MoSCoW is best for a release deadline, and our constraint is ongoing capacity. Let me propose a different lens.”*
- The framework chosen to confirm a preference. If the product lead proposes RICE specifically because it’s the lens that makes their preferred item win, name it carefully: *“Let’s pick the lens that fits the decision, not the lens that fits the answer.”*

Phase 3: Silent scoring (15 min)

Hand out the scoring sheet, paper or a private tab, with every item listed. Each person scores every item against the framework’s axes. Silent. No discussion.

For MoSCoW, each person assigns a letter per item. For RICE, each person writes four numbers per item and the sheet computes. For Cost of Delay, each person writes a delay-cost and a size. For value-vs-effort, each person places a dot on a 2×2.

Clarifying questions about what an item is go on a side-sheet and get answered only when everyone has finished scoring. The facilitator holds silence actively, stepping in when someone starts to narrate, and politely redirecting side conversations.

What to watch for:

- Anchoring. If one person finishes early and starts commenting, the still-scoring participants anchor to their voice. Cut it off: *"Let's finish scoring first."*
- Effort inflation. When everything lands at a 5, 8, or 13 on a Fibonacci-like sizing scale (1, 2, 3, 5, 8, 13) someone is avoiding the hard estimate. In the reveal, call it: *"These nine items all have the same effort. Do we believe that?"*
- Silent disagreement. A participant scores the founder's pet feature low and doesn't want to reveal it. The session is designed to surface exactly this; protect the silent phase so their score gets recorded before group pressure kicks in.

Phase 4: Reveal and discuss (25 min)

Collect the scores onto a single sheet (the wall, a projected spreadsheet, a shared doc). Show each item with every participant's score and the aggregate.

Start with agreement. Items where the scores are tight: read them out, confirm they're in the top or bottom as expected, move on. Agreement at the edges is a gift; the session doesn't need to revisit it.

Then walk the disagreements. For each item where the scoring distribution is wide (scores spanning from low to high) open a short conversation:

"This one has scores from 2 to 9. I'd like the person with the 9 and the person with the 2 to explain what they were seeing."

The goal is not to reach consensus. The goal is to expose the hidden context behind the extremes. Often one person has information the other didn't (a conversation with a customer, a bug they saw, a piece of architecture they know about). Making that visible is the value of the reveal.

What to watch for:

- Re-scoring under pressure. Someone revises their score after hearing someone else's argument. Sometimes fine, sometimes conformity. Ask: *"Did your model actually change, or did you just hear a louder voice?"*
- The HiPPO (Highest Paid Person's Opinion). The highest-paid person's opinion anchors the room. Surface it: *"I notice we're all moving toward [name]'s score. Anyone still holding a different view?"*
- Information-hoarding. Someone scored high because of a fact only they know. Make them share it; that's the point of the reveal.

Phase 5: Argue the fringes (20 min)

The top and the bottom of the ordered list are where prioritisation earns its keep. Walk the top three and the bottom three out loud.

For the top:

"These three are the items we're committing to first. Does anyone in the room think one of these is wrong?"

Silence is acceptance. A raised hand is valuable data. The person who raises their hand usually has a reason you haven't yet heard.

For the bottom:

"These three are the items we're not doing this cycle. Is that right? Is there anything here we'd regret not doing?"

The bottom-up check catches items whose scores were low because of information gaps. The CS voice often saves an item here that the scoring missed.

What to watch for:

- Everything's a top item. The room refuses to commit to the bottom. Usually a constraint problem: nobody believes you actually can't do all of it. Return to the constraint card.
- Quiet disagreement in front of the founder. The top looks wrong but nobody's saying. Ask directly, by name: *"You haven't said much this round. What do you think of the top three?"*
- A rescued bottom item. Someone argues an item out of the bottom. Good, but it displaces something. Force the exchange: *"If we pull this up, what drops?"*

Phase 6: Land the order and commit (10 min)

The product lead reads the final order aloud. Everyone confirms. The facilitator photographs the list, the scoring sheet, and the wall.

"Product lead: is this the order you're committing to? Everyone: does anyone here have a reason to ring the alarm bell on this order before we commit?"

State commitments out loud, by name. *"I'll turn the top three into backlog items by Monday. Engineering lead: you'll take the top item into planning Tuesday."*

End the session on a commitment, not a summary.

See [Prioritisation: What Changes First](#) for the Greenbox team's first prioritisation session after the backlog outgrew the whiteboard, including the moment Tom's 2 against Maya's 9 on the same item turns out to be the most valuable conversation of the week.

What Can Go Wrong

The missing constraint. The room starts prioritising without agreeing what they're optimising against.

Recovery: Stop the session. *"We can't prioritise without a constraint. Ten minutes to name one."* *Stop if:* The room can't agree on a constraint in fifteen minutes. The prioritisation isn't the problem; something upstream is.

HiPPO drift. Scores migrate toward whoever is most senior in the room. *Recovery:* Ask the quietest person to speak first on the next item. Keep anonymous scoring sheets if you can; don't read names against scores until the distribution is on the wall. *Stop if:* The senior person keeps interrupting silent scoring. Pause the session and have a private word. If they won't hold silence, they should not be in the session.

Effort inflation. Every item is an 8 or a 13. The room is avoiding the estimate. *Recovery:* *"Let's look at the three smallest items. If these aren't smaller than the others, the t-shirt sizing isn't helping. What would a 2 look like?"* *Stop if:* The team genuinely doesn't know. The backlog isn't refined enough for prioritisation; run a refinement session first.

Framework quibble. The room argues about the framework rather than the items. *Recovery:* “We can use any framework well or any framework badly. Let’s score with the one on the card.” *Stop if:* The quibble doesn’t end after five minutes. The facilitator misjudged the framework; pick a different one and restart.

Quiet disagreement. The top three look wrong to someone who isn’t saying. *Recovery:* Name them and ask directly. “You scored this item a 3. The room’s score is 8. What were you seeing?” *Stop if:* They won’t speak even when asked. There’s a trust problem; that’s not fixable in the session.

Ambiguous win condition. Halfway through, the room realises the constraint was ambiguous. *Recovery:* Pause scoring. Clarify the constraint. Re-score items whose scores depended on the ambiguity. *Stop if:* The clarification reveals the constraint was wrong, not ambiguous. End the session, name the real constraint, reconvene.

The session that doesn’t survive contact with Monday. The order is produced, and the backlog is reordered differently two days later without revisiting the session. *Recovery:* Pin the constraint card and the photographed order in the team’s tracker. When someone proposes a re-order, point at the card and ask which item it displaces. *Stop if:* The product lead won’t defend the order. The session was theatre; the real decision is being made elsewhere.

The deeper failure modes worth naming:

- Silent scoring gets skipped and the session devolves into a debate
- The constraint is fake and the scores are theatre
- The team confuses “we picked a framework” with “we decided”. The framework is input, the decision is the product lead’s
- The top three get committed but the bottom three aren’t killed, so the backlog grows rather than shrinks

The costs of running it honestly:

- 4-7 people × 60-90 minutes, plus the product lead’s half-day of prep
- Emotional cost when someone’s favourite item lands at the bottom
- Recurring cost. Prioritisation is a living conversation; one session does not last a quarter
- Political cost when the chosen lens produces an order the founder didn’t expect

Next Steps

The session ends; the work begins.

Same day, the facilitator:

- Photographs the scoring wall, the constraint card, and the final ordered list.
- Transcribes the list into the team’s tracker, with the constraint and the framework noted at the top.
- Writes a short summary for anyone not in the room: here’s the constraint, here’s the lens, here’s the top three, here’s the bottom three, here’s why.

This week, the product lead:

This is where the pattern earns its cost.

- Turn the top three into real backlog items within 48 hours. If the top three don't move from prioritisation cards into the tracker quickly, the momentum dies and the session is forgotten.
- Kill or park the bottom three explicitly. A stated "not now" with a date is more valuable than a silent drift. Users and stakeholders get told: *"We looked at this, we're not doing it this cycle, here's when we'll revisit."*
- Close the scoring loop. When an item turns out to have been much bigger or smaller than scored, bring it back to the team. Calibrating the scoring is how the next session gets faster.
- Defend the order in the face of new requests. The hardest week-after task: when a stakeholder proposes new urgent work, they have to displace something on the list, not just appear alongside it.

Ongoing, the team:

- Re-runs prioritisation on a cadence that matches the constraint's horizon. A fortnight for tactical; a quarter for portfolio; a year for strategic.
- Tracks which items that scored low were secretly important (the bottom-rescue items). A pattern there points to a missing scoring axis.
- Keeps the constraint card pinned. When someone asks why an item was deprioritised, point to the card.

Variants

Tactical fortnight (default). A team's top-of-backlog conversation, 10 to 20 candidate items, 60 to 90 minutes. Constraint is usually capacity for the next fortnight. Value-vs-effort or RICE fits most teams; Cost of Delay when the work is mixed. This is what most teams need, and the rest of this post describes it.

Portfolio quarter. A leadership conversation across multiple teams, 20 to 30 items spanning features, platform work, and migrations. Cost of Delay / WSJF is the natural lens because the items have genuinely different urgency profiles. Run on a quarterly cadence; the output buckets cleanly into Now / Next / Later for stakeholder-facing roadmaps.

Strategic year. A founder-and-senior-team conversation about which initiatives to fund. Smaller item count, larger items, MoSCoW or Now / Next / Later as the framing. Run annually or when the strategic shape genuinely shifts; the Business Model Canvas often feeds this version.

Release MoSCoW. A scope conversation tied to a fixed deadline. MoSCoW is the right lens here and only here: when the output is a yes/no decision per item against a launch date. Time-boxed launches, MVP scoping, regulated rollouts.

Remote. A shared doc or board with one row per item, private tabs for silent scoring (each person fills their column without seeing others), then reveal by un-hiding the columns. Slightly slower than in-person (the rhythm of placing scores on a wall is faster) but the silence is actually easier to protect remotely. Use one shared cursor during reveal so the facilitator drives the discussion.

Backlog-rescue marathon. Two or three sessions back-to-back over a half-day when the backlog has spiralled and a single 90-minute slot won't clear it. Frame each session against a different constraint (this fortnight, next quarter, the year) and let the same items get scored under each lens. The cross-lens disagreements are where the real strategic conversations live.

About this playbook

This playbook is part of *The Workshop*, a reference series of facilitator playbooks published at barkingiguana.com. The canonical, up-to-date version lives at barkingiguana.com/writing/the-workshop-prioritisation/.

These posts are LLM-aided. Backbone, original writing, and structure by Craig. Research and editing by Craig + LLM. Proof-reading by Craig.